

Self-organization, Emergence and the Architecture of Complexity

Francis HEYLIGHEN*

Transdisciplinary Research Group, Free University of Brussels, Pleinlaan 2, B-1050 Brussels, Belgium

ABSTRACT. It is argued that the problems of emergence and the architecture of complexity can be solved by analysing the self-organizing evolution of complex systems. A generalized, distributed variation-selection model is proposed, in which internal and external aspects of selection and variation are contrasted. "Relational closure" is introduced as an internal selection criterion. A possible application of the theory in the form of a pattern directed computer system for supporting complex problem-solving is sketched.

RESUME. La thèse est proposée selon laquelle les problèmes d'émergence et de l'architecture de la complexité peuvent être résolues en analysant l'évolution auto-organisatrice des systèmes complexes. Un modèle généralisé, distribué de variation-et-sélection est proposé, et les aspects internes et externes de la variation et de la sélection sont mis en contraste. La "fermeture relationnelle" est introduite comme critère de sélection interne. Une application potentielle de la théorie, un système informatique aidant à résoudre les problèmes complexes, est envisagée.

1. Emergence and self-organization

Emergence is a classical concept in systems theory, where it denotes the principle that the global properties defining higher order systems or "wholes" (e.g. boundaries, organization, control, ...) can in general not be reduced to the properties of the lower order subsystems or "parts". Such irreducible properties are called emergent. Until now there is no satisfactory theory explaining what characterizes emergent properties or what are the conditions for their existence. In this paper I propose to look at this question not from the traditional static viewpoint but from a dynamic, evolutionary viewpoint, replacing the question "How can a property *be* emergent?" by "How can a property *become* emergent? (i.e. how can it *emerge*?)". This should also lead us to answer the question "Where do 'wholes' or 'systems' come from?"

A promising approach to the problem of dynamical emergence is provided by the recently developed models of self-organization. Self-organization may be defined as a spontaneous (i.e. not steered or directed by an external system) process of organization, i.e. of the development of an organized structure. The spontaneous creation of an "organized whole" out of a "disordered" collection of interacting parts,

* Senior Research Assistant NFWO (Belgian National Fund for Scientific Research)

as witnessed in self-organizing systems in physics, chemistry, biology, sociology ..., is a basic part of dynamical emergence.

However, another essential characteristic of emergence as it is understood in systems theory is its hierarchical or multi-level nature: an emergent whole at one level is merely a component of an emergent system at the next higher level. Until now, the most popular paradigms used for explaining self-organization (e.g. attractors, synergetics, catastrophes, ...) are characterized by a mere two-level structure: the "microscopic" level where a multitude of building blocks or elements (e.g. molecules, individual organisms, ...) interact, and the "macroscopic" level where these interactions lead to certain global patterns of organization (e.g. a dissipative structure or a crystalline symmetry). The resulting systems as studied through these paradigms (e.g. a crystal, a regular pattern of fluid rolls in the Bénard phenomenon, or a trail of ants carrying food back to the nest) are usually so simple in structure that it is not necessary to use a specifically systemic approach for understanding them.

2. Emergence and the architecture of complexity

Realistically complex systems (e.g. organisms, societies, ecologies...), however, are characterized by a multi-level structure. A classic explanation for this hierarchical "architecture" of complex systems was given by Simon (1962). His argument is based on a variation-and-selection view of natural (and artificial) evolution: elements are connected and combined by natural interactions (or, equivalently, by the trials of a problem-solver), thus creating a variety of assemblies. Of these assemblies only those will "survive" which are sufficiently stable, the other assemblies will fall apart before they can undergo any further evolution. The stable assemblies, forming "naturally selected wholes", can then again function as building blocks, to be combined into higher order assemblies, and so the process can repeat itself at ever higher levels, forming a set of hierarchically structured complexes.

Simon then uses this model in order to show why multi-level systems are more probable to emerge than two-level systems of comparable complexity: in a two-level system *all the components* must "fall into place" at once, otherwise the assembly will be unstable and fall apart before the missing components are added by the natural variation mechanisms. In a multi-level system, on the other hand, it suffices that *small subsets of components* would "fall into place" forming stable subassemblies ("modules"), which can then again be recursively combined in small sets forming higher level modules. Clearly, the smaller the set of elements which must fall into place, the higher the probability that this will happen by random combination.

However, Simon acknowledges that there are exceptions to this rule that non-hierarchical complex systems are highly improbable: for example, most polymers are formed by a very simple linear, two-level assembly of a large number of molecules. One of the important contributions of present-day self-organization models is that they can explain the emergence of such non-modular, two-level systems, which have nevertheless a very large number of elements. Such processes are usually characterized by non-linear, autocatalytic mechanisms, whereby the presence of a small stable assembly (whose emergence is quite probable) enhances the probability that other elements would join the assembly, thus making it grow and become even more stable. (in the formulation of Haken (1983): a stable mode "slaves" the

remaining unstable modes). No intermediate levels of modules are needed in such a process with positive feedback. The emergent stable configuration can be thought of as an "attractor" exerting a force on the configurations in its neighbourhood, so that the configurations which are close enough to the attractor will automatically move closer and closer towards this stable configuration.

It is clear that both the hierarchical model of Simon and the "non-linear" models of self-organization only describe part of the features of emergence. A real complex system (e.g. the human body) has as well hierarchical, multi-level aspects (e.g. the organelle being a subsystem of the cell, being a subsystem of the organ, being a subsystem of...) as non-linear, two-level aspects (e.g. the system of blood vessels as a coordinated closed circuit consisting of billions of blood cells). However, in general, there is not just one global hierarchy or non-linear organization, but a multitude of inextricably entwined suborganizations and subsystems. If we wish to understand the architecture of such complexity, we will need a more general, integrating theory of emergence and self-organization. The present text will propose some basic principles on which such a theory could be founded.

3. A generalized variation-and-selection dynamics

The theory of natural selection as it is used for describing biological evolution can be simply generalized to any kind of systemic evolution. It suffices to consider a system undergoing variation and an environment exerting a "selective pressure" on the system: only those configurations of the system will maintain (or grow) which are "fit" or adapted to the environment. The evolving system can be likened to a *problem-solver*, generating possible solutions by trial (variation) to a problem posed by the environment: how to be optimally adapted? The problem arises as soon as adaptation is not optimal, i.e. the system is not perfectly stable or invariant with respect to the environment. The larger the instability, the more serious the problem, and the more variation the system must undergo before it reaches a new equilibrium.

Like all good problem-solvers know, it does not suffice to blindly try out possibilities, in the hope that accidentally one of them would prove to be the optimal solution: the chances that this would succeed are very small. You can enhance your chances by looking for intermediate steps, i.e. relatively easy-to-find problem states or configurations, which are no final solutions but which are somehow "closer" to the goal than the configuration you started with. This is what also happens during natural selection, as exemplified by the models of Simon and the self-organization theorists discussed above.

It should be noted here that there is no essential difference between "*deterministic*" and "*indeterministic*" processes of variation or trial: you can either systematically search through all states according to a given rule (e.g. depth-first), or try out states in a certain domain at random. Although in the first case you can predict which state will be studied at step N of the process, you cannot predict whether this state will be a solution (i.e. will be adapted or stable). The property of "being a solution" is *emergent*, i.e. it cannot be explained or predicted at the level where the search is carried out. It only appears when the state is actually reached (cfr. Heylighen, 1989c). Therefore any process of variation or search is "blind", although it is not necessarily random (cfr. Campbell, 1974). The "problem of

determinism" can be left for metaphysics, and is irrelevant for the study of concrete, adaptive systems.

The difference between (traditional) problem-solving and evolution, however, is that in the latter case there is no final solution. A system is never optimally adapted to an environment since the process of evolution of the system will itself change the environment so that a new adaptation is needed, and so on. Every "goal" of the process can be seen as a subgoal of another, further away goal, which is again a subgoal of ... Hence the "subgoals" become the important features of the process. They correspond to what we have called stable assemblies or emerging systems.

Another difference between the above, simple model and more complex evolutions is that evolution is in general *parallel* or *distributed*: there is not just one system and its environment, there is a multitude of systems evolving simultaneously, partially autonomously, partially in interaction. This "network" structure of evolutionary processes entails that no absolute distinction can be made between internal and external, i.e. between system and environment. What is "system" for one process is "environment" for another one.

This means that natural selection can no longer be interpreted simply as selection by the environment. One way to evade this problem is to look at the whole of systems evolving in parallel as just one global system (e.g. an ecology). In that case natural selection means that the variation of the global system leads to globally stable configurations. External selection has now been replaced by internal selection: the internal structure of the system must be stable for the system to survive; we do not need to look at its adaptation to an external environment. Of course, in practice it is impossible to study an absolutely global system (i.e. the universe), and so each practical system will have as well an aspect of *internal selection* (intrinsic stability) and an aspect of *external selection* (adaptation).

The thesis I want to advance is that until now too little attention has been paid to the internal (i.e. self-organizing) aspect of selection. Indeed any external selection can be reduced to internal selection by considering a larger, more global system. Adaptation is then reduced to the existence of a stable relation between one subsystem (the original system) and another subsystem (the original environment).

For example, an "external factor" with a selective influence on plants is the amount of carbon dioxide in the atmosphere, which the plant needs for survival. This factor can be seen as a "boundary condition", as an environmental constraint to which the self-organizing processes in the plant must adapt. Yet from a more global viewpoint, the carbon dioxide is not a given, external condition but a product of the adaptation of another type of systems : animals and bacteria. These systems depend on another selective factor: the oxygen in the atmosphere, which is produced by the plants from the carbon dioxide. Hence we see that the double, parallel adaptation process of plants and animals can be seen as the internal self-organization of the ecology as a whole, leading to the selection of a stable cycle in which carbon dioxide is transformed to oxygen and back to carbon dioxide.

The duality between internal and external selection can also be applied to variation. *Internal variation* may be defined as a process in which an inner part of a system is changed. An example is mutation in which one gene within a chromosome is changed. *External variation* means that the relation between the system and its environment is changed, i.e. that the system is coupled to different external systems. An example is chromosome recombination during sexual reproduction in which one chromosome of a pair is coupled to a chromosome of another pair from another

organism. Again, what is external from one viewpoint (the viewpoint of the individual chromosome) may be internal from another one (the viewpoint of the gene pool of the species within which sexual reproduction is possible), and vice-versa.

4. Relational closure as an internal stability criterion

As shown above, a variation-selection dynamics which is general, in the sense that it can be applied to systems of any type in any environment, should be based on internal stability as a universal selection criterion. Indeed, no external selection criterion can be universal since it depends on the environment chosen.

If we wish to define such an internal stability criterion we must first determine what should be stable or invariant. Clearly we cannot demand that the system as a whole, with all its components, properties and relations with other systems, should remain invariant. In that case no evolution or change whatsoever would be possible. The essential feature which must be maintained is the *identity* of the system. This identity can be defined as that which *distinguishes* the system from its background or environment. This allows the emergence of systems whose state may change, but whose identity is maintained.

An example of such systems are autopoietic systems, as described by Varela (1979). They are characterized by a changing structure ("state") and an invariant organization ("identity"). This invariant identity is the result of what Varela (1979) calls "organizational closure": the network of processes constituting the system is organized in such a way that it produces itself. This concept is rather difficult, and the typical examples (biological organisms) are very complex, so that it is not very clear how a universal, mathematical theory of self-organization could be based on it. Therefore, I propose to introduce an extension of this concept: "relational closure", which can be analysed mathematically and conceptually in a very simple way.

In mathematical systems theories "closure" is understood as the invariance of a set under an algebra (e.g. a group) of transformations (cfr. Ashby, 1964). The problem with this definition is that there is an a priori separation between algebras (representing systems), individual transformations (representing subsystems or subprocesses) and sets. In a general self-organizing complex it is not clear which features should be modelled as algebras and which as transformations. However, any system can be modelled as a relation between an input set and an output set (Mesarovic & Takahara, 1975). By going to the level of relational algebra it is not even necessary to introduce sets: relations are defined by the way they can be composed or coupled. (in set terminology: a relation R can be composed with R' if the output set (domain) of R corresponds with the input set (codomain) of R'). Hence we could model a system as a relation which is itself formed by a network of relations (subsystems) which are either coupled (can be composed) or not. (Indeed, any sequence of coupled relations can be reduced to a single relation by composition, any set of parallel (i.e. not coupled) relations can be reduced to a single relation by set-theoretic union).

The relational closure of a system can then be defined as the *internal invariance of a distinction* (or distinction system) defining the system (Heylighen, 1989a,b). Internal invariance means that the distinction is mapped upon itself during the subsequent application of the transformations (i.e. subrelations) inherent in the system (remember that a relation between elements is mathematically equivalent to a

function mapping elements or subsets onto subsets). The closure of an assembly of subsystems then signifies that the subsystems are connected (through their input and output) in such a way that some overall order, structure or organization may be distinguished, which is invariant under the dynamical processes generated by the interaction between the subsystems.

An example of such an organization is that of a transformation group. Whatever the dynamics of the system, activating one transformation after another one, the resulting changes will always remain "within the system". The concept of a "closed" system is more general than that of a group, however. A transformation group is merely a very clear example of the generic concept, because it combines several, more elementary types of closures: *transitive* or recursive closure (the internality of composition in a group), *cyclical* closure (the existence of inverse transformations), and *surjective* (many-to-one) and *inverse surjective* (one-to-many) closure (the bijectivity of group transformations) (Heylighen, 1989a).

The different combinations of these 4 basic closures define a wealth of closed structures, which are weaker than that of a group. For example, the combination of transitivity and cyclicity defines an equivalence relation. The associated invariant distinction is that between separate equivalence classes in a partition. The combination of transitivity and non-cyclicity defines a partial order. The corresponding distinction is that between elements which are "higher" or "lower" with respect to the ordering. The addition of surjectivity to this combination results in a tree structure or hierarchy; the further addition of inverse surjectivity defines a linear order. Complete cyclicity (i.e. symmetry) and non-transitivity define an orthogonality relation. The addition of surjective or inverse surjective closure leads to "orthogonally closed" classes. These form the base for orthocomplemented lattices and hence for Boolean algebras. Spencer-Brown (1969) has shown that these algebras are isomorphic to what he defines as "distinction algebras".

I do not want to exhaust all possible combinations of elementary closures and their corresponding mathematical structures and distinctions here. I just want to argue (without at this stage being able to prove it formally) that all fundamental types of abstract order or organization which can be found in emerging systems, such as hierarchies, symmetries, periodicities, cycles, partitions, ..., can be generated by the recursive combination of extremely simple "closure" operations. For example, a limit cycle, which is the prototype of an attractor in self-organization theories, corresponds obviously to a cyclically closed dynamical system. A crystal symmetry, on the other hand, is characterized by a group of symmetry transformations. Such elementary closures may be recursively combined forming higher-order closed systems, which themselves may form the building blocks of still higher order closed systems, ..., and so on, without limit.

Let us apply the closure concept to the distributed evolution mechanism sketched above. The external variation can be viewed as the creation and destruction of couplings between systems. Some of these temporary couplings may be such that the resulting system is closed, in the sense defined above. In that case the assembly will gain some form of internal stability, so that it survives long enough to be used as a building block or template for subsequent variation processes, leading to new couplings and eventually to new closures. Because of the closure a new distinction, characterizing the closed assembly, is created. This distinction cannot be reduced to just a combination of the more elementary distinctions characterizing the subsystems

of the assembly, since it "integrates" and replaces these distinctions on a higher level. It thus truly corresponds to an "emerging property".

It must be remarked here that such distinctions arising from closure are more general than just boundaries separating the inside of a system from its outside. For example the distinction between "higher" and "lower" in an order relation (e.g. "later" and "earlier" in a time ordering) is a structural characteristic which cannot be reduced to a relation between "inside" and "outside". This allows us to get a better understanding of the architecture of the complexity resulting from the evolution process as sketched. If the only emerging distinctions were boundary distinctions (represented e.g. by equivalence classes) then we would get a purely hierarchic architecture, with boundaries of subsystems enclosed within the boundaries of their supersystems, and so on. The emergence of cyclical distinctions, on the other hand, may lead to heterarchical, non-linear architectures. Group distinctions then may lead to flat, repetitive structures, ...

Since the same element or subsystem may be part of different "closures" or higher order systems, it is clear that these different organizations or architectures cannot be neatly separated out. In a sense the overall architecture of the complex is a *superposition* of all these partial organizations. This concept may be illustrated by considering a representation of the human body as it is used in some anatomy books: the model consists of a sequence of superposed transparent sheets, each sheet containing a map of a particular subsystem of the body: the skeleton, the circulatory system, the nervous system, the digestive system, the muscular system, ... It is clearly impossible to picture all these subsystems in one two-dimensional representation, since they are braided together in an extremely complex way. Indeed, the same element—organ or cell—will in general belong to several subsystems. For example, the heart is a muscle and the central part of the circulatory system; moreover it contains nerves and the lymphatic vessels flow between its cells. Yet each subsystem performs an individual function, and all of its elements collaborate in a coherent relationally "closed" manner, so that it can be unambiguously distinguished from the other subsystems.

The "superposition" of closed suborganizations can also be compared with the superposition of states in quantum mechanics. Observations of a quantum mechanical system only result in eigenstates of the observable property which was measured. Such eigenstates may be considered "closed" (i.e. invariant) with respect to the operation of measurement of the corresponding property. Yet a general state of the system will not be an eigenstate of a particular property, but it can always be represented as a "superposition" of such eigenstates.

5. Operationalizing the theory

The concepts and principles introduced above should not remain purely theoretical speculations. With the advent of the new information technology complex, qualitative mechanisms can now be implemented and tested on computer in a relatively simple way. A general programming paradigm, *pattern directed systems*, is emerging, which is directly applicable to the present type of approach. A pattern directed system consist of a collection of modules or rules, which respond to messages ("conditions") characterized by a specific pattern (i.e. a set of variables or input channels structured in a specific way) by sending out new messages

("actions"), dependent on the information received. The system is intrinsically parallel since different modules can respond simultaneously to the same (or different) message(s), but it is possible to simulate such mechanisms on sequential machines. Examples of pattern directed systems are : production systems, classifier systems, object-oriented systems, and logical or relational programming.

In our approach the modules can be likened to (sub)systems, the messages to their input and output. Two modules can be said to be (temporarily) coupled if the output message of the one is accepted as input by the other one. The general problem with pattern directed systems is to specify the control structure, i.e. the set of rules which determines which module can send or accept messages to or from which other module. The generalized variation-selection dynamics in combination with the closure concept may provide an answer to this problem.

The dynamics controlling the flow of messages must depend on two selection criteria: the external "problem", to be specified by the user, and the internal closure of collections of coupled rules, leading to the self-organization and emergence of complex subsystems within the pattern directed system. In order to be effective the system should also have a variation mechanism. In order to start the problem-solving (= evolution) process, there must be an original variety of modules. This can be provided by the user, who could try to express the initial knowledge he has about the problem domain in the form of "if ... then ..." modules. Of course, this initial variety can always be expanded by the user during the problem-solving process: there is a continuous interaction between the computer system and the user, who plays the role of the external environment. Another source of variety can be provided by the computer system itself, which generates variations of the existing modules by internal changes or by combinations with different, external modules. Until now, typical problem-solving programs (working according to the generate-and-test mechanism) only use internal variation, i.e. the state of the system is changed by replacing some of its intrinsic properties. However, we have shown that external variation is a more interesting process in the sense that it can give rise to the emergence of higher-order systems through closure.

An example of an existing pattern directed system evolving through variation-selection is formed by "classifier systems" (Wilson, 1987). Here the selection is basically external, but the variation is partially internal ("mutation" of classifiers), partially of a mixed type ("recombination" of classifiers, in which part of one module (=classifier) is recombined with part of another module). There is no explicit closure mechanism. Moreover, the information contained in a module is fixed, so that there is no explicit mechanism for emergence, although complex "assemblies" of modules might implicitly develop.

Let us conclude by sketching how a pattern directed implementation of the present theory of emergence and evolution might be applied to "real world" problems. The main idea would be to design a generic computer support system for solving complex problems (Heylighen, 1989b). A problem, as said, can be defined as a situation of non-optimal or non-satisfactory adaptation. The problem does not need to be well-structured (i.e. have an explicit goal, initial state and domain), it suffices that the actor experiencing the problem be capable of distinguishing satisfactory solutions from non-satisfactory ones, i.e. that he be able to carry out a selection between possibilities offered to him. The task of the support-system would then be to provide the user with potential solutions, with a relatively high probability of success.

Therefore the system must possess some intelligence, i.e. use the available (though usually incomplete) knowledge in an efficient way by integrating the pieces of knowledge in stable, adaptive systems or complexes, and adapt itself rapidly to new input from the user. Moreover, the proposed potential (or partial) solutions should be meaningful to the user, i.e. easily recognizable as satisfactory or not. Therefore, the organization of the proposed system should be transparent and motivated. This demands an advanced interface for representing complex information. Such an interface may be provided with the aid of so-called "hypermedia" (e.g. HyperCard on the Apple Macintosh), i.e. the combination of multiple media (text, graphics, sound, programming, animation...) in a non-sequential, but easily accessible, network format. Furthermore the system should continuously offer advice and explanations regarding the possible evolutions of the problem-solving process. In that way the system should be able to support a non-expert user in formulating, structuring and solving problems of any domain.

References

- Ashby W. R. (1964) : *Introduction to Cybernetics*, (Methuen, London).
- Campbell D.T. (1974) : *Evolutionary Epistemology*, in : The Philosophy of Karl Popper, Schilpp P.A. (ed.), (Open Court Publishing, La Salle, Illinois), p. 413.
- Haken H. (1983) : *Synergetics : Non-Equilibrium Phase Transitions and Self-Organization in Physics, Chemistry and Biology* (3rd ed.), (Springer, Berlin).
- Heylighen F. (1989a): *Building a Science of Complexity*, Proceedings of the 1988 Annual Conference of the Cybernetics Society (London).
- Heylighen F. (1989b) : *Coping with Complexity: concepts and principles for a support system*, Preceedings of the conference "Support, Society and Culture", R. Glanville & G. de Zeeuw (eds.) (OOC, University of Amsterdam).
- Heylighen F. (1989c): *Causality as Distinction Conservation: a theory of predictability, reversibility and time order*, submitted to Synthese.
- Mesarovic M.D. & Takahara Y. (1975) : *General Systems Theory : Mathematical Foundations*, (Academic Press, New York).
- Simon H.A. (1962): *The Architecture of Complexity*, Proceedings of the American Philosophical Society 106, p. 467-482, reprinted in:
- Simon H.A. (1981) : *The Sciences of the Articial* (2nd ed.), (MIT Press, Cambridge MA).
- Spencer Brown G. (1969) : *Laws of Form*, (Allen & Unwin, London).
- Varela F.J. (1979) : *Principles of Biological Autonomy*, (North Holland, New York).
- Wilson S.W. (1987): *Classifier Systems and the Animat Problem*, Machine Learning 2, p. 199-228.